

# On optimal direction gibbs sampling

J. Andrés Christen<sup>(a)\*</sup>, Colin Fox<sup>(b)</sup>,  
Diego Andrés Pérez-Ruiz<sup>(a)</sup> and Mario Santana-Cibrian<sup>(a)</sup>

(a) Centro de Investigación en Matemáticas (CIMAT),  
A.P. 402, Guanajuato, Gto. 36000, Mexico. *jac@cimat.mx*

(b) Colin Fox, Department of Physics, University of Otago,  
New Zealand. *fox@physics.otago.ac.nz*

May 2012

## Abstract

Generalized Gibbs kernels are those that may take any direction not necessarily bounded each axis along the parameters of the objective function. We study how to optimally choose such directions in a Directional, random scan, Gibbs sampler setting. The optimal direction is chosen by minimizing to the mutual information (Kullback-Leibler divergence) of two steps of the MCMC for a truncated Normal objective function. The result is generalized to be used when a Multivariate Normal (local) approximation is available for the objective function. Three Gibbs direction distributions are tested in highly skewed non-normal objective functions.

Keywords: MCMC; Bayesian inference; simulation.

## 1 Introduction

Let  $\pi(\mathbf{x})$  be the objective distribution of interest. Let  $\mathbf{X} \in \mathbb{R}^n$  be a random variable with density  $\pi(\mathbf{x})$ . The (univariate) Gibbs sampler is a MCMC sampling algorithm that simulates systematically or randomly from the conditional distributions

$$f_{X_i|\mathbf{X}_{-i}}(x_i|\mathbf{x}_{-i}) \propto \pi(\mathbf{x}), \quad (1)$$

where the notation

$$\mathbf{v}_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$$

represents the  $(n - 1)$ -dimension vector created by deleting the  $i$ -th entry from the  $n$ -dimension vector  $\mathbf{v}$ . (1) above represent univariate distributions and are the

---

\*Corresponding author.

conditional distributions along the axis for the base chosen to represent  $\mathbf{X}$ , the so called full conditional distributions.

Using (1) a Markov chain  $\mathbf{X}^{(1)}, \mathbf{X}^{(2)}, \dots$  is created considering the transition kernel

$$K_i(\mathbf{x}^{(t)}, \mathbf{x}^{(t+1)}) = f_{X_i|\mathbf{X}_{-i}}(x_i^{(t+1)}|\mathbf{x}_{-i}^{(t)})1(\mathbf{x}_{-i}^{(t+1)} = \mathbf{x}_{-i}^{(t)}).$$

That is, the  $i$ -th kernel changes only the  $i$ -th coordinate by simulating from the full conditional distribution  $f_{X_i|\mathbf{X}_{-i}}(\cdot|\mathbf{x}_{-i}^{(t)})$ . In random scan Gibbs, a complete transition kernel is defined as

$$K(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^n w_i K_i(\mathbf{x}, \mathbf{y}),$$

for some weights  $\sum_{i=1}^n w_i = 1$ ;  $w_i \geq 0$ .

The direction Gibbs sampler generalizes this idea by choosing a direction  $\mathbf{e} \in \mathbb{R}^n$ ,  $\|\mathbf{e}\| = 1$ , and sampling from the conditional distribution along such direction. This can be written as

$$\mathbf{X}^{(t+1)} = \mathbf{x}^{(t)} + r\mathbf{e},$$

where the length  $r \in \mathbb{R}$  has distribution proportional to  $\pi(\mathbf{x}^{(t)} + r\mathbf{e})$ . In any case, it can be seen that the transition kernel has detailed balance with respect to  $\pi$ , and by assuring  $\pi$ -irreducibility the Markov chain has  $\pi$  as ergodic distribution.

The natural question to ask is how to choose  $\mathbf{e}$  to optimize the convergence (mixing) of the Markov chain? Indeed, once irreducibility is assured, any chain will have the correct ergodic distribution ( $\pi$ ) but performance will depend on how dependent  $\mathbf{X}^{(t+1)}$  and  $\mathbf{X}^{(t)}$  are. In this context a convenient, although less known, dependence measure to use is the *mutual information*  $I$  between random variables  $X$  and  $Y$ , which measures the Kullback-Leibler divergence between the joint model  $f_{X,Y}$  and the independent alternative  $f_X f_Y$ , that is,

$$I(\mathbf{Y}, \mathbf{X}) = \int \int f_{\mathbf{Y},\mathbf{X}}(\mathbf{y}, \mathbf{x}) \log \frac{f_{\mathbf{Y},\mathbf{X}}(\mathbf{y}, \mathbf{x})}{f_{\mathbf{Y}}(\mathbf{y})f_{\mathbf{X}}(\mathbf{x})} d\mathbf{x}d\mathbf{y}.$$

In our case, this translates to  $\mathbf{X} = \mathbf{X}^{(t)}$  and  $\mathbf{Y} = \mathbf{X}^{(t+1)}$ . Assuming that  $\mathbf{X} \sim \pi$ , we see that  $f_{\mathbf{Y}}(\mathbf{y}) = \pi(\mathbf{y})$  and  $f_{\mathbf{Y},\mathbf{X}}(\mathbf{y}, \mathbf{x}) = \pi(\mathbf{x})K(\mathbf{x}, \mathbf{y})$ . Therefore, the mutual information above can be calculated as

$$I(\mathbf{Y}, \mathbf{X}) = \int \int \pi(\mathbf{x})K(\mathbf{x}, \mathbf{y}) \log \frac{K(\mathbf{x}, \mathbf{y})}{\pi(\mathbf{y})} d\mathbf{y}d\mathbf{x}. \quad (2)$$

The idea would be to choose directions for which  $I(\mathbf{X}^{(t+1)}, \mathbf{X}^{(t)})$  is minimized. Since  $I$  is a Kullback-Leibler divergence it is well defined,  $I \geq 0$  and  $I = 0$  if and only if  $\mathbf{X}^{(t+1)}$  and  $\mathbf{X}^{(t)}$  are independent ( $f_{X,Y} = f_X f_Y$  a.s.). Finding directions  $\mathbf{e}$  for which  $I$  is minimum will provide our optimization criterion, to obtaining better suited direction Gibbs samplers.

## 2 The Normal case

In many applications, a Multivariate Normal approximation to the posterior is possible. We will try to use such Normal approximation to produce an improved Gibbs samplers. In this section we will assume  $\pi$  to be Normal, a case that lends itself to calculate  $I(\mathbf{X}^{(t+1)}, \mathbf{X}^{(t)})$ . For the moment, we will consider that the dimension is sufficiently low so as to have the eigen decomposition of the precision matrix available.

We therefore assume that  $\pi$  is a multi-variate Normal with *precision* ( $n \times n$ ) matrix  $\mathbf{A}$  (the inverse of the variance-covariance matrix) and mean column vector  $\boldsymbol{\mu}$ . As explained above, given a direction  $\mathbf{e} \in \mathbb{R}^n$ ,  $\|\mathbf{e}\| = 1$ , the chain moves from  $\mathbf{X}^{(t)} = \mathbf{x}$  to

$$\mathbf{Y} = \mathbf{X}^{(t+1)} = \mathbf{x} + r\mathbf{e},$$

where the length  $r \in \mathbb{R}$  has distribution  $g$  proportional to  $\pi(\mathbf{x} + r\mathbf{e})$ . That is

$$g(r) \propto \exp \left\{ -\frac{1}{2}(\mathbf{v} + r\mathbf{e})' \mathbf{A}(\mathbf{v} + r\mathbf{e}) \right\},$$

where  $\mathbf{v} = \mathbf{x} - \boldsymbol{\mu}$ . After some algebra we see that  $r \sim N \left( -\frac{\mathbf{e}' \mathbf{A} \mathbf{v}}{\mathbf{e}' \mathbf{A} \mathbf{e}}, \mathbf{e}' \mathbf{A} \mathbf{e} \right)$ , where  $\mathbf{e}' \mathbf{A} \mathbf{e}$  is the *precision* (by setting  $\mathbf{e} = \mathbf{e}_i$ , the  $i$ -th standard bases vector, one obtains

$$Y_i \sim N \left( \mu_i - a_{ii}^{-1}(\mathbf{v}' \mathbf{v} - (x_i - \mu_i)^2), a_{ii} \right),$$

which is the full conditional distribution for a Multivariate Normal distribution for entry  $i$ , thus returning to the usual Gibbs sampler). From this we see that the transition kernel corresponding to direction  $\mathbf{e}$  is

$$K_{\mathbf{e}}(\mathbf{x}, \mathbf{y}) = \left( \frac{\mathbf{e}' \mathbf{A} \mathbf{e}}{2\pi} \right)^{\frac{1}{2}} \exp \left\{ -\frac{\mathbf{e}' \mathbf{A} \mathbf{e}}{2} \left( \mathbf{e}'(\mathbf{y} - \mathbf{x}) + \frac{\mathbf{e}' \mathbf{A} \mathbf{v}}{\mathbf{e}' \mathbf{A} \mathbf{e}} \right)^2 \right\} 1(\mathbf{y} = \mathbf{x} + \mathbf{e}'(\mathbf{y} - \mathbf{x})\mathbf{e})$$

(note that  $\mathbf{y} - \mathbf{x} = r\mathbf{e}$  and since  $\mathbf{e}' \mathbf{e} = 1$ ,  $r = \mathbf{e}'(\mathbf{y} - \mathbf{x})$ ; therefore  $\mathbf{y}$  is restricted to the line  $\mathbf{y} = \mathbf{x} + \mathbf{e}'(\mathbf{y} - \mathbf{x})\mathbf{e}$ ).

We calculate  $I_{\mathbf{e}}(\mathbf{X}^{(t+1)}, \mathbf{X}^{(t)})$ , the mutual information of the Gibbs sampler given direction  $\mathbf{e}$ , as in 2. Note that

$$\log \frac{K_{\mathbf{e}}(\mathbf{x}, \mathbf{y})}{\pi(\mathbf{y})} = C + \frac{1}{2} \log \mathbf{e}' \mathbf{A} \mathbf{e} - \frac{1}{2} (Q_1(\mathbf{e}, \mathbf{x}, \mathbf{y}) - Q_2(\mathbf{y})),$$

where

$$\begin{aligned} C &= \frac{n-1}{2} \log 2\pi - \frac{1}{2} \log |\mathbf{A}|, \\ Q_1(\mathbf{e}, \mathbf{x}, \mathbf{y}) &= \mathbf{e}' \mathbf{A} \mathbf{e} \left( \mathbf{e}'(\mathbf{y} - \mathbf{x}) + \frac{\mathbf{e}' \mathbf{A} \mathbf{v}}{\mathbf{e}' \mathbf{A} \mathbf{e}} \right)^2, \\ Q_2(\mathbf{y}) &= (\mathbf{y} - \boldsymbol{\mu})' \mathbf{A} (\mathbf{y} - \boldsymbol{\mu}). \end{aligned}$$

From this we see that

$$\int \log \frac{K_e(\mathbf{x}, \mathbf{y})}{\pi(\mathbf{y})} K_e(\mathbf{x}, \mathbf{y}) d\mathbf{y} = C - \frac{1}{2} + \frac{1}{2} \log \mathbf{e}' \mathbf{A} \mathbf{e} + \frac{1}{2} \int Q_2(\mathbf{y}) K_e(\mathbf{x}, \mathbf{y}) d\mathbf{y}$$

since  $\int Q_1(\mathbf{e}, \mathbf{x}, \mathbf{y}) K_e(\mathbf{x}, \mathbf{y}) d\mathbf{y} = 1$ . The integral  $\int Q_2(\mathbf{y}) K_e(\mathbf{x}, \mathbf{y}) d\mathbf{y}$  may be calculated by transforming back to  $r$  since

$$\int Q_2(\mathbf{y}) K_e(\mathbf{x}, \mathbf{y}) d\mathbf{y} = \int (r\mathbf{e} - \mathbf{v})' \mathbf{A} (r\mathbf{e} - \mathbf{v}) g_e(r) dr.$$

After some algebra one sees that

$$\int Q_2(\mathbf{y}) K_e(\mathbf{x}, \mathbf{y}) d\mathbf{y} = 1 - \frac{\mathbf{v}' \mathbf{A} \mathbf{e} \mathbf{e}' \mathbf{A} \mathbf{v}}{\mathbf{e}' \mathbf{A} \mathbf{e}} + \mathbf{v}' \mathbf{A} \mathbf{v}.$$

Therefore

$$\int \log \frac{K_e(\mathbf{x}, \mathbf{y})}{\pi(\mathbf{y})} K_e(\mathbf{x}, \mathbf{y}) d\mathbf{y} = C + \frac{1}{2} \log \mathbf{e}' \mathbf{A} \mathbf{e} - \frac{1}{2} \frac{\mathbf{v}' \mathbf{A} \mathbf{e} \mathbf{e}' \mathbf{A} \mathbf{v}}{\mathbf{e}' \mathbf{A} \mathbf{e}} + \mathbf{v}' \mathbf{A} \mathbf{v}.$$

We need now to integrate with respect to  $\pi(d\mathbf{x})$ . We note that  $\int \mathbf{v}' \mathbf{A} \mathbf{v} \pi(\mathbf{x}) d\mathbf{x} = n$ . Moreover, the expected value of a quadratic form is

$$E(\mathbf{z}' \mathbf{R} \mathbf{z}) = \text{tr}(\mathbf{R} \mathbf{\Sigma}) + \boldsymbol{\mu}' \mathbf{R} \boldsymbol{\mu},$$

where  $\boldsymbol{\mu}$  and  $\mathbf{\Sigma}$  are the mean vector and the variance-covariance matrix of  $\mathbf{z}$ . Letting  $\mathbf{R} = \frac{\mathbf{A} \mathbf{e} \mathbf{e}' \mathbf{A}}{\mathbf{e}' \mathbf{A} \mathbf{e}}$  and since  $E(\mathbf{v}) = \mathbf{0}$  we obtain

$$\begin{aligned} E\left(\mathbf{v}' \frac{\mathbf{A} \mathbf{e} \mathbf{e}' \mathbf{A}}{\mathbf{e}' \mathbf{A} \mathbf{e}} \mathbf{v}\right) &= \frac{1}{\mathbf{e}' \mathbf{A} \mathbf{e}} \text{tr}(\mathbf{A} \mathbf{e} \mathbf{e}' \mathbf{A} \mathbf{A}^{-1}) \\ &= \frac{1}{\mathbf{e}' \mathbf{A} \mathbf{e}} \text{tr}(\mathbf{A} \mathbf{e} \mathbf{e}') \\ &= \frac{1}{\mathbf{e}' \mathbf{A} \mathbf{e}} \text{tr}(\mathbf{e} \mathbf{A} \mathbf{e}') \\ &= 1. \end{aligned}$$

Therefore

$$\begin{aligned} I_e(\mathbf{X}^{(t+1)}, \mathbf{X}^{(t)}) &= C + n - \frac{1}{2} + \frac{1}{2} \log \mathbf{e}' \mathbf{A} \mathbf{e} \\ &= C_1 + \frac{1}{2} \log \mathbf{e}' \mathbf{A} \mathbf{e}, \end{aligned} \tag{3}$$

where  $C_1 = C + n - \frac{1}{2}$ .

## 2.1 Choosing a set of directions

We need now a distribution  $h$  for directions to be chosen to generate an *irreducible* Gibbs sampler. That is, according to (3) the best direction is that that minimizes  $C_1 + \frac{1}{2} \log \mathbf{e}' \mathbf{A} \mathbf{e}$ , but we cannot simply choose the best direction: the resulting Gibbs will not be irreducible and clearly we will not be sampling from  $\pi$ . The chain must be  $\pi$ -irreducible in order that the chain has as ergodic distribution  $\pi$ . Indeed, if directions have distribution  $h(\mathbf{e})$ , and this distribution has the whole sphere  $\mathbb{S}^n$  as support, then the resulting Markov chain  $K(\mathbf{x}, \mathbf{y}) = \int K_e(\mathbf{x}, \mathbf{y}) h(\mathbf{e}) d\mathbf{e}$  is irreducible.

Kaufman and Smith (1994) argue that an optimal direction distribution is

$$h(\mathbf{e}) \propto \sup_{\mathbf{x} \in \mathcal{X}, r \in \mathbb{R}} \left\{ \int \pi(\mathbf{x} + \tau \mathbf{e}) d\tau \frac{|r|^{n-1}}{\pi(\mathbf{x} + r \mathbf{e})} \right\},$$

as far as optimizing the (geometric) rate of convergence of the resulting Gibbs sampling. However, this only applies for a bounded support  $\mathcal{X}$  for  $\pi$ . Little else has been said regarding the optimal directions for generalized Gibbs samplers. We cannot control in general the term  $\frac{|r|^{n-1}}{\pi(\mathbf{x} + r \mathbf{e})}$  for unbounded support. However, this suggests choosing a direction

$$h^*(\mathbf{e}) \propto \sup_{\mathbf{x} \in \mathbb{R}^n} \left\{ \int \pi(\mathbf{x} + \tau \mathbf{e}) d\tau \right\}.$$

For the normal case, it is not difficult to see that

$$\int \pi(\mathbf{x} + \tau \mathbf{e}) d\tau \leq \frac{(2\pi)^{-\frac{n-1}{2}} |\mathbf{A}|^{\frac{1}{2}}}{\sqrt{\mathbf{e}' \mathbf{A} \mathbf{e}}} \exp \left\{ \frac{1}{2} \frac{(\mathbf{e}' \mathbf{A} \mathbf{v})^2}{\mathbf{e}' \mathbf{A} \mathbf{e}} \right\}.$$

From this we take  $h^*(\mathbf{e}) \propto (\mathbf{e}' \mathbf{A} \mathbf{e})^{-\frac{1}{2}}$ . On the other hand, one can minimize  $I_e(\mathbf{X}^{(t+1)}, \mathbf{X}^{(t)})$  by maximizing  $\exp\{-I_e(\mathbf{X}^{(t+1)}, \mathbf{X}^{(t)})\}$ . Then, choosing

$$h^*(\mathbf{e}) \propto (\mathbf{e}' \mathbf{A} \mathbf{e})^{-\frac{1}{2}}$$

will naturally choose directions with low  $I_e(\mathbf{X}^{(t+1)}, \mathbf{X}^{(t)})$ , as can be seen from (3).

To sample from  $h^*(\mathbf{e})$ , it is also easy to see that  $h^*(\mathbf{e}) \propto \int \pi(\boldsymbol{\mu} + \tau \mathbf{e}) d\tau \propto (\mathbf{e}' \mathbf{A} \mathbf{e})^{-\frac{1}{2}}$  (the maximum is achieved when  $\mathbf{x} = \boldsymbol{\mu}$ ). If we simulate  $\mathbf{e}_u$  from a Multivariate Normal with precision matrix  $\mathbf{A}$  and centered at the origin,  $\pi_0$ , and take  $\mathbf{e} = \frac{\mathbf{e}_u}{\|\mathbf{e}_u\|}$ , it is clear that  $\mathbf{e}$  will have density  $\propto \int \pi_0(\tau \mathbf{e}) d\tau \propto (\mathbf{e}' \mathbf{A} \mathbf{e})^{-\frac{1}{2}}$ . That is,  $\mathbf{e} \sim h^*$ , and this density has the whole sphere  $\mathbb{S}^n$  as its support and this results in an ergodic chain.

## 3 Non-normal objective function

So far nothing has been achieved (since for sampling from a multivariate normal several samples from basically the same MN distribution are needed!). But indeed, the aim is to produce an efficient sampler when a (local) Multivariate Normal approximation exists. Note from the previous section that our Optimal Direction

Gibbs algorithm only requires knowledge of the precision matrix  $\mathbf{A}$ , not of the mean  $\boldsymbol{\mu}$  (to sample the direction from  $h^*$ ). That is, only an (local) approximation to the the log of the objective function will be needed.

### 3.1 Truncated Normal objective function

Sometimes the quantities we want to determine are different from the ones which we are able to measure. If the data measured depends on the quantities we want, then the data contains some information about those quantities. Inverse problems occurs when observed data  $\mathbf{y}$  depend on unknowns  $\mathbf{x}$  via a measurement process, and we want to recover  $\mathbf{x}$  from  $\mathbf{y}$ . In linear inverse problems, the relationship between  $\mathbf{y}$  and  $\mathbf{x}$  is given by

$$\mathbf{B}\mathbf{x} = \mathbf{y},$$

where  $\mathbf{B}$  is a linear transformation. The classical inverse problem is to invert the function  $\mathbf{B}$  to obtain unknowns  $\mathbf{x}$  in terms of data. Suppose that  $\mathbf{B}_{m \times n} \mathbf{x}_{n \times 1} = \mathbf{y}_{m \times 1}$  is a linear inverse problem where the matrix  $\mathbf{B}$  is known. Suppose that  $\mathbf{y}|\mathbf{x} \sim N_m(\mathbf{B}\mathbf{x}, \mathbf{T})$  with  $\mathbf{T} = \text{diag}(\frac{1}{\sigma_1^2}, \dots, \frac{1}{\sigma_m^2})$ , and suppose  $\mathbf{x}$  has a Truncate Multivariate Normal Distribution with known mean vector  $\boldsymbol{\mu}$  and precision matrix  $\mathbf{A}$ , where  $x_i > 0 \forall i$ . We are interested in find the posterior distribution of  $\mathbf{x}|\mathbf{y}$ . After some algebra, we see that

$$f(\mathbf{x}|\mathbf{y}) \propto \exp \left\{ -\frac{1}{2}((\mathbf{x} - \boldsymbol{\mu}^*)' \mathbf{A}^* (\mathbf{x} - \boldsymbol{\mu}^*)) \right\} \\ \cdot 1(x_i > 0)$$

is a Truncate Multivariate Normal Distribution, with precision matrix

$$\mathbf{A}^* = (\mathbf{A} + \mathbf{B}'(m\mathbf{T})\mathbf{B}),$$

and mean vector

$$\boldsymbol{\mu}^* = (\mathbf{A} + \mathbf{B}'(m\mathbf{T})\mathbf{B})^{-1}(\boldsymbol{\mu}'\mathbf{A} + \bar{\mathbf{y}}'(m\mathbf{T})\mathbf{B}).$$

We take a Multivariate Normal distribution  $\pi$  with precision  $(n \times n)$  matrix  $\mathbf{A}$  and mean vector  $\boldsymbol{\mu} = \left(\sqrt{\frac{1}{n}}, \dots, \sqrt{\frac{1}{n}}\right)$ , but truncated the support to  $x_i \geq 0$  (all entries are positive, our objective function is  $\propto \pi(\mathbf{x})$ ). A random direction Gibbs sampling is considered with direction distribution  $h^*$  as in Section 2.1. We set the initial point  $\mathbf{X}^{(0)} = \boldsymbol{\mu}$  and the burn in is not needed. The full conditionals are the same as in Section 2 but bound by the positivity constraint.

We consider several dimensions  $n = 2, 3, 5, 10, 15, 20$ . In each case, we compare with a standard random scan (alone the base axis) Gibbs sampling, a pure random direction Gibbs and directions simulated from  $h^*$ . Using the QR decomposition of a  $n \times n$  matrix of uniform random entries one obtains a random orthonormal matrix  $\mathbf{P}$ , which will represent the orthonormal base of eigenvalues. The eigenvectors of  $\mathbf{A}$ ,  $\lambda_1, \dots, \lambda_n$ , represent the precisions on each eigenvalue direction and are set to

$\lambda_i = \sigma_i^{-2}$ . The standard deviations in each principal (eigen) direction are set to

$$\sigma_i = i^{\frac{-\alpha}{n}}, \quad (4)$$

$\alpha \geq 0$ , and  $\mathbf{A} = \mathbf{P}'\mathbf{\Lambda}\mathbf{P}$ , where  $\mathbf{\Lambda} = \text{diag}(\lambda_i)$ . These represent decreasing standard deviations and are increasingly contrasting as  $\alpha$  increases. More contrasting standard deviations result in further correlated distributions. We consider  $\alpha = 0, 5, 10, 20$ .

For each combination of  $n$  and  $\alpha$  we calculate the Integrated Autocorrelation Time of the resulting chains. Results are shown in Figure 1. Note that the random scan and random direction Gibbs worsen in performance as the  $\sigma_i$ 's are more contrasting, our Optimal Direction Gibbs (ODG) remains with an  $\text{IAT}/n \leq 12$ . Note that in this general case, the ODG is a random walk; the resulting IAT's are quite close to the theoretical optimum for an optimal scaled random walk, as explained in (Roberts and Rosenthal, 2001).

### 3.2 More general objective functions

Suppose then that  $\pi(\mathbf{x})$  is the objective distribution of interest with support  $\mathcal{X} \subset \mathbb{R}^n$ . Let  $\mathbf{X} \in \mathbb{R}^n$  be a random variable with density  $\pi(\mathbf{x})$ . Suppose also that for each  $\mathbf{x} \in \mathcal{X}$  we have the Hessian  $H(\mathbf{x})$  for  $-\log \pi(\cdot)$  evaluated at  $\mathbf{x}$ . By completing the squares in the second order Taylor approximation to  $-\log \pi(\cdot)$ , we assume sufficient smoothness such that

$$-\log \pi(\mathbf{y}) \approx C(\mathbf{x}) + \frac{1}{2}(\mathbf{y} - \boldsymbol{\mu}(\mathbf{x}))' H(\mathbf{x})(\mathbf{y} - \boldsymbol{\mu}(\mathbf{x})),$$

where  $\boldsymbol{\mu}(\mathbf{x}) = \mathbf{x} - H^{-1}(\mathbf{x})\nabla(\mathbf{x})$  ( $\nabla(\mathbf{x})$  is the gradient of  $-\log \pi(\cdot)$  evaluated at  $\mathbf{x}$ ). That is, we have a local MN approximation to  $\pi$  with mean vector  $\boldsymbol{\mu}(\mathbf{x})$  and precision matrix  $H(\mathbf{x})$ .


The proposed algorithm runs as follows. At  $\mathbf{X}^{(t)} = \mathbf{x}$  propose jumping to

$$\mathbf{Y} = \mathbf{x} + r\mathbf{e} \quad (5)$$


where the length  $r \in \mathbb{R}$  has distribution  $r \sim N\left(-\frac{\mathbf{e}'H(\mathbf{x})\mathbf{v}}{\mathbf{e}'H(\mathbf{x})\mathbf{e}}, \mathbf{e}'H(\mathbf{x})\mathbf{e}\right)$  and  $\mathbf{v} = \mathbf{x} - \boldsymbol{\mu}(\mathbf{x}) = \mathbf{x} - \mathbf{x} + H^{-1}(\mathbf{x})\nabla(\mathbf{x})$ . This translates to

$$\begin{aligned} r &\sim N\left(-\frac{\mathbf{e}'H(\mathbf{x})H^{-1}(\mathbf{x})\nabla(\mathbf{x})}{\mathbf{e}'H(\mathbf{x})\mathbf{e}}, \mathbf{e}'H(\mathbf{x})\mathbf{e}\right) \\ &= N\left(-\frac{\mathbf{e}'\nabla(\mathbf{x})}{\mathbf{e}'H(\mathbf{x})\mathbf{e}}, \mathbf{e}'H(\mathbf{x})\mathbf{e}\right). \end{aligned}$$


As usual, the proposal  $\mathbf{Y}$  will be accepted with a Metropolis-Hastings probability. In the case that  $\pi$  is a MN as in Section 2 it is clear that  $\nabla(\mathbf{x}) = \mathbf{A}(\mathbf{x} - \boldsymbol{\mu})$ ,  $H(\mathbf{x}) = \mathbf{A}$ , therefore  $\mathbf{v} = \mathbf{A}^{-1}\mathbf{A}(\mathbf{x} - \boldsymbol{\mu})$  and we are back with the same sampler.




(a)



(b)



(c)



(d)

Figure 1: Integrated autocorrelation times divided by the dimension of the MN objective function. Direction gibbs for random scan (blue), random direction (green) and direction according to  $h^*$  (red). The standard deviations for the MN are set according to (4) for (a)  $\alpha = 0$  (independent normals) (b)  $\alpha = 5$ , (c)  $\alpha = 10$  and (d)  $\alpha = 20$ .



The Metropolis-Hasting ratio would be

$$R = \frac{\pi(\mathbf{y}) q(\mathbf{x}|\mathbf{y})}{\pi(\mathbf{x}) q(\mathbf{y}|\mathbf{x})},$$

where  $q(\cdot|\mathbf{x}) = h^*(\mathbf{e})q_e(\cdot|\mathbf{x})$  is the density of  $\mathbf{Y}$  in (5). In Section 2  $q_e(\cdot|\mathbf{x})$  is the conditional distribution over the direction  $\mathbf{e}$ , given  $\mathbf{x}$  (gibbs kernel), and cancels out with  $\pi(\mathbf{x})$  ( $h^*(\mathbf{e})$  would cancel out with  $h^*(-\mathbf{e})$  since the Hessian is constant;  $H(\mathbf{x}) = \mathbf{A}$ ). Let  $g(\mathbf{x}) = \log \pi(\mathbf{x})$ , then

$$\log R = (g(\mathbf{y}) + \log q(\mathbf{x}|\mathbf{y})) - (g(\mathbf{x}) + \log q(\mathbf{y}|\mathbf{x})).$$

Let

$$\psi_{\mathbf{x}}(\mathbf{e}, r) = g(\mathbf{x}) + \log q(\mathbf{x} + r\mathbf{e}|\mathbf{x}).$$

Note that  $q(\mathbf{x} + r\mathbf{e}|\mathbf{x}) = h^*(\mathbf{e})q_e(\mathbf{x} + r\mathbf{e}|\mathbf{x})$  and therefore

$$\begin{aligned} q(\mathbf{x} + r\mathbf{e}|\mathbf{x}) &= K_{H(\mathbf{x})}(\mathbf{e}'H(\mathbf{x})\mathbf{e})^{-\frac{1}{2}} \frac{(\mathbf{e}'H(\mathbf{x})\mathbf{e})^{\frac{1}{2}}}{\sqrt{2\pi}} \\ &\cdot \exp \left\{ -\frac{\mathbf{e}'H(\mathbf{x})\mathbf{e}}{2} \left( r + \frac{\mathbf{e}'\nabla(\mathbf{x})}{\mathbf{e}'H(\mathbf{x})\mathbf{e}} \right)^2 \right\}, \end{aligned}$$

and

$$\psi_{\mathbf{x}}(\mathbf{e}, r) = g(\mathbf{x}) + \log K_{H(\mathbf{x})} - \frac{\mathbf{e}'H(\mathbf{x})\mathbf{e}}{2} \left( r + \frac{\mathbf{e}'\nabla(\mathbf{x})}{\mathbf{e}'H(\mathbf{x})\mathbf{e}} \right)^2,$$

where

$$K_{H(\mathbf{x})}^{-1} = \int_{\|\mathbf{e}\|=1} (\mathbf{e}'H(\mathbf{x})\mathbf{e})^{-\frac{1}{2}} d\mathbf{e}$$

is the normalizing constant of  $h^*$  that now depends on  $\mathbf{x}$ . We then have

$$R(\mathbf{x}, \mathbf{e}, r) = \exp\{\psi_{\mathbf{x}+r\mathbf{e}}(-\mathbf{e}, r) - \psi_{\mathbf{x}}(\mathbf{e}, r)\}. \quad (6)$$

Therefore, the probability of accepting a jump from  $\mathbf{x}$  to  $\mathbf{y} = \mathbf{x} + r\mathbf{e}$  should be  $\min\{1, R(\mathbf{x}, \mathbf{e}, r)\}$ . However, it is not possible to obtain  $K_{H(\mathbf{x})}$  analytically. For dimension  $n = 2$  we calculate  $K_{H(\mathbf{x})}$  numerically, only for comparison purposes.

An alternative distribution of directions would be the following. Assuming that  $\mathbf{A} = H(\mathbf{x})$  is positive defined, we can use its eigenvalues and eigenvectors. We will take the directions  $\mathbf{e}$  as the eigenvectors of matrix  $H(\mathbf{x})$ , so  $\mathbf{e} \in \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\}$ . The  $i$ -th direction will be selected with probability proportional to  $\lambda_i^{-1}$ , where  $\lambda_i$  is the eigenvalue corresponding to the  $i$ -th eigenvector,  $i = 1, 2, \dots, n$ . Then  $h_1(\mathbf{e}_i) = (k\lambda_i)^{-1}$ , where  $k = \sum_{i=1}^n \lambda_i^{-1}$ . By doing this the distribution of  $\mathbf{e}$  becomes discrete. Moreover, we avoid the problem we have with the normalization constant  $K_{H(\mathbf{e})}$  and the implementation can be done for large values of dimension  $n$ .

It is easy to see that direction  $\mathbf{e}_1$ , corresponding to the lowest eigenvalue  $\lambda_n$  of

	$\alpha$	$\Sigma$	IAT	Acceptance rate
a)	$(-1.0, -1.0)$	$(1.0, 0.5, 0.5, 1.0)$	4.039870	0.8712
b)	$(-0.5, 5.0)$	$(1.0, 0.9, 0.9, 1.0)$	6.944900	0.7693
c)	$(-5.0, 5.0)$	$(1.0, 0.9, 0.9, 1.0)$	3.909896	0.8485
d)	$(-10.0, -10.0)$	$(1.0, 0.5, 0.5, 1.0)$	7.629253	0.6892

Table 1: Acceptance rate and IAT for simulations obtained with the Optimal Direction Gibbs algorithm using distribution  $h*$ .

$H(\mathbf{x})$ , is optimal indeed. Note that

$$\begin{aligned} \min_{\|\mathbf{e}\|=1} I_{\mathbf{e}}(\mathbf{X}^{(t+1)}, \mathbf{X}^{(t)}) &= \min_{\|\mathbf{e}\|=1} \left\{ C + \frac{1}{2} \log(\mathbf{e}' H(\mathbf{x}) \mathbf{e}) \right\} \\ &= C + \frac{1}{2} \log \left( \min_{\|\mathbf{e}\|=1} \{\mathbf{e}' H(\mathbf{x}) \mathbf{e}\} \right) = C + \frac{1}{2} \log \lambda_1. \end{aligned}$$

The minimum is reached when  $\mathbf{e} = \mathbf{e}_1$ , the eigenvector associated to  $\lambda_1$ .

### 3.3 Example: The skew-normal distribution

To test the performance of our algorithm we will use as objective distribution a variant of the skew-normal distribution. We choose this distribution due its simplicity and versatility, plus the fact that it allows us to calculate analytically the gradient and the Hessian of  $-\log \pi$ .

The multivariate skew-normal probability density function is given by

$$f(\mathbf{y}) = 2\phi_n(\mathbf{y} - \boldsymbol{\xi}; \Sigma) \Phi(\boldsymbol{\alpha}'(\mathbf{y} - \boldsymbol{\xi})),$$

where  $\phi_n(\cdot; \Sigma)$  is the multivariate normal probability density function with mean vector  $\mathbf{0}$  and variance-covariance matrix  $\Sigma$ ,  $\Phi(\cdot)$  is the standard univariate normal cumulative distribution function,  $\boldsymbol{\xi}$  is a localization parameter and  $\boldsymbol{\alpha}$  is a shape parameter (Azzalini and Adelchi, 2005).

$\Phi(\cdot)$  is called *perturbation* function and may be the cumulative distribution function of any univariate symmetric distribution around 0. In our case, we are going to use the Logistic cumulative distribution function with mean 0 and scale parameter  $s = \frac{\sqrt{3}}{\pi}$ , which is

$$G(x) = \frac{1}{1 + \exp\{-\pi x/\sqrt{3}\}}.$$

Then, the density function of our objective distribution is given by

$$\pi(\mathbf{x}) = 2 \left( \frac{|A|}{(2\pi)^n} \right)^{1/2} \exp \left\{ -\frac{1}{2} \mathbf{x}' \mathbf{A} \mathbf{x} \right\} \frac{1}{1 + \exp\{-\pi \boldsymbol{\alpha}' \mathbf{x} / \sqrt{3}\}},$$




Figure 2: Simulations obtained with the Optimal Direction Gibbs algorithm using distribution  $h^*$ .

where  $\mathbf{A} = \Sigma^{-1}$ . In order to implement the algorithm we need to calculate the gradient and Hessian of  $-\log \pi(\mathbf{x})$ . It is easy to see that

$$\nabla(\mathbf{x}) = \mathbf{A}\mathbf{x} - (1 - G(\boldsymbol{\alpha}'\mathbf{x}))\boldsymbol{\alpha},$$

and

$$H(\mathbf{x}) = \mathbf{A} + g(\boldsymbol{\alpha}'\mathbf{x})\boldsymbol{\alpha}\boldsymbol{\alpha}',$$

where  $g(\cdot)$  is the Logistic probability density function.

To test the algorithm we consider the case  $n = 2$ . Figure 2 shows the contour plots of the objective distribution for different values of  $\Sigma$  and  $\boldsymbol{\alpha}$ . It also shows the simulations obtained using the algorithm with distribution  $h^*$ , after 10,000 iterations. We set the initial point  $\mathbf{X}^{(0)} = \mathbf{0}$  and the burn in is not needed. Table 1 shows the acceptance rate and IAT for each example. Note that the value of the IAT is small and the acceptance rate is large in all cases. Both measures indicate that the performance of the algorithm is quite good. Is important to mention that a conventional Gibbs sampler cannot be used for this example since the full conditional distributions do not have any close form. Our scheme has not that kind of difficulties.

Unfortunately, the implementation is not as simple for dimension  $n \geq 3$  since we need to calculate the normalization constant  $K_{H(\mathbf{x})}$  numerically.

In Subsection 3.2 we proposed an alternative direction distribution  $h_1$  to avoid calculating  $K_{H(\mathbf{x})}$ . We test the algorithm using this distribution on the same examples. Figure 3 shows some very interesting results. Simulations in examples a) and b) are well distributed over the region of interest. However, in examples c) and d) simulations are concentrated in a specific region. This results from the fact that the eigenvalues of  $H(\mathbf{x})$  are very contrasting, one is much larger than the other, and consequently, simulations are obtained in similar directions constantly.

This suggests that the scheme works well in simple cases, but it becomes inefficient against more skewed distributions. However, we can still modify  $h_1$ . We take the directions  $\mathbf{e}$  as eigenvectors of  $H(\mathbf{x})$ , but the probability of choosing  $\mathbf{e}_i$  will now be proportional to  $(\lambda_i)^{-b}$ , where  $b \sim \text{Beta}(1, 9)$ , this is  $h_2(\mathbf{e}_i) \propto (\lambda_i)^{-b}$ . By doing this we allow the chance that probabilities become more balanced in regions where there is a very dominant eigenvalue.

Figure 4 shows the behaviour of the resulting algorithm. Examples a) and b) show a good performance as before. On the other hand, examples c) and d), which we identified as difficult above, seem to have been corrected. Table 2 shows the acceptance rate and IAT. Here we see good results. Acceptance rates remain high and do not represent a problem. The IAT is slightly higher than in Table 1, except by example c); however, the difference is not quite significant.

Direction distribution  $h_2$  seems a good alternative for our optimal direction distribution since it shows good performance and the IAT values are very similar to those obtained with our original scheme. The parameters of the Beta distribution were selected intuitively. In fact, those parameters allows us to change distribution  $h_2$  in order to get better results.

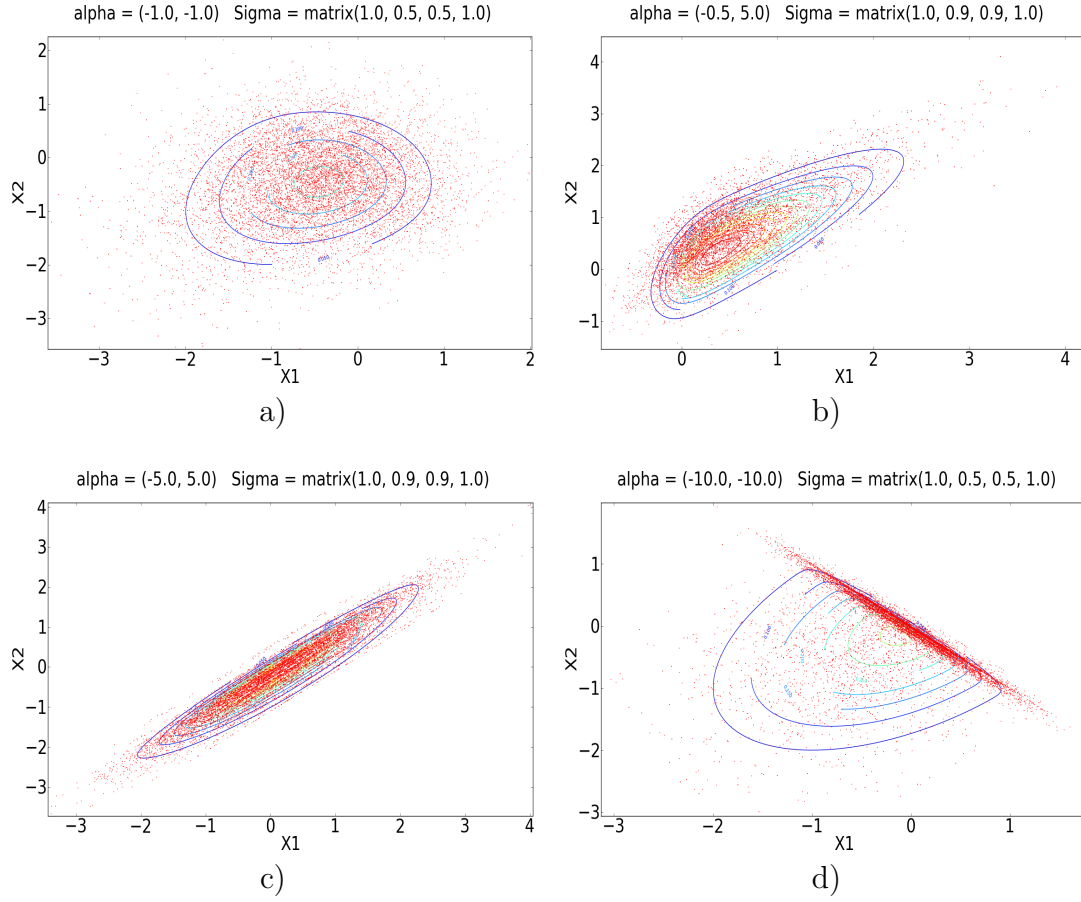


Figure 3: Simulations obtained with the Optimal Direction Gibbs algorithm using distribution  $h_1$ .

	$\alpha$	$\Sigma$	IAT	Acceptance rate
a)	$(-1.0, -1.0)$	$(1.0, 0.5, 0.5, 1.0)$	4.497301	0.8793
b)	$(-0.5, 5.0)$	$(1.0, 0.9, 0.9, 1.0)$	7.844630	0.7534
c)	$(-5.0, 5.0)$	$(1.0, 0.9, 0.9, 1.0)$	2.705416	0.8809
d)	$(-10.0, -10.0)$	$(1.0, 0.5, 0.5, 1.0)$	8.821470	0.7327

Table 2: Acceptance rate and IAT for simulations obtained with the Optimal Direction Gibbs algorithm using distribution  $h_2$ .

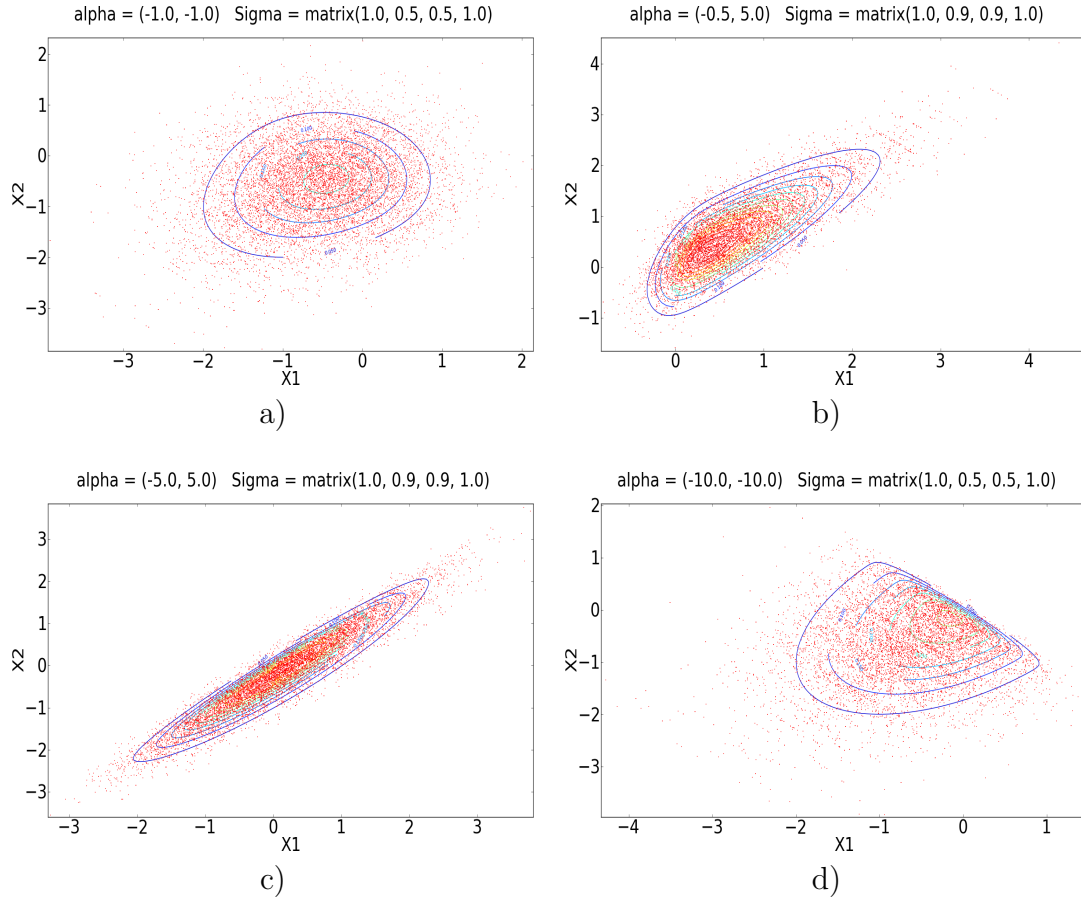


Figure 4: Simulations obtained with the Optimal Direction Gibbs algorithm using distribution  $h_2$ .

## 4 Discussion

Our Optimal Direction Gibbs sampler presents interesting characteristics in examples where either a conventional gibbs sampler is impossible to implement or other MCMC methods (eg. a Random Walk Metropolis-Hastings) would be very difficult to tune. The truncated normal example is of great relevance in the field of inverse problems and we have also worked with strongly skewed distributions with contrasting scales, with promising results in all cases.

## 5 Acknowledgements

DAPR and MSC thank CONACyT for a MSc scholarship since part of this work was conducted while finishing their MSc studies at CIMAT. Part of JAC work was founded by CONACyT grant 128477.

## References

- Azzalini, Adelchi (2005) The skew-normal distribution and related multivariate families. *Scandinavian Journal of Statistics* 2:159–188
- Kaufman D, Smith R (1994) Direction choice for accelerated convergence in hit-and-run sampling. *Operation Research* 46(1):84–95
- Roberts GO, Rosenthal JS (2001) Optimal scaling for various metropolis-hastings algorithms. *Statistical Science* 16(4):351–367